Modeling of Transitional Channel Flow Using Balanced Proper Orthogonal Decomposition

Miloš Ilak* and Clarence W. Rowley†

Department of Mechanical and Aerospace Engineering,

Princeton University, Princeton, NJ 08544, USA

(Dated: February 5, 2008)

arXiv:0707.4112v2 [math.OC] 21 Dec 2007

Abstract

We study reduced-order models of three-dimensional perturbations in linearized channel flow using balanced proper orthogonal decomposition (BPOD). The models are obtained from three-dimensional simulations in physical space as opposed to the traditional single-wavenumber approach, and are therefore better able to capture the effects of localized disturbances or localized actuators. In order to assess the performance of the models, we consider the impulse response and frequency response, and variation of the Reynolds number as a model parameter. We show that the BPOD procedure yields models that capture the transient growth well at a low order, whereas standard POD does not capture the growth unless a considerably larger number of modes is included, and even then can be inaccurate. In the case of a localized actuator, we show that POD modes which are not energetically significant can be very important for capturing the energy growth. In addition, a comparison of the subspaces resulting from the two methods suggests that the use of a non-orthogonal projection with adjoint modes is most likely the main reason for the superior performance of BPOD. We also demonstrate that for single-wavenumber perturbations, low-order BPOD models reproduce the dominant eigenvalues of the full system better than POD models of the same order. These features indicate that the simple, yet accurate BPOD models are a good candidate for developing model-based controllers for channel flow.

PACS numbers: 47.27.nd,47.27.Cn,47.85.L-

*Electronic address: milak@princeton.edu

†Electronic address: cwrowley@princeton.edu

I. INTRODUCTION

Many techniques for developing practical controllers for fluids require models of the system that are both tractable and that describe accurately the flow physics for the given flow regime. One of the problems of great interest in flow control is drag reduction in shear flows. Drag increases drastically as flow transitions from laminar to turbulent, making turbulence suppression or inhibition of transition to turbulence promising strategies for drag reduction. Both open-loop and closed-loop strategies have been used in control of channel flow. Recently, Min et al [22] have achieved sub-laminar drag through open-loop control using a traveling wave actuation on the channel walls. Lee et al [20] have reported drag reduction using closed-loop controllers developed for the linearized flow, while Joshi et al [15] and Högberg et al [12] have demonstrated significant reduction in the energy of the perturbations to laminar flow using closed-loop controllers. Despite these recent successes, the adequate modeling of the fluid flow and the actuators in a framework useful for practical control design is still a challenge. In order to use model-based control strategies, one needs an accurate description of the system dynamics and the actuation, from a model of sufficiently low order to allow practical implementation. The goal of this paper is to explore improved techniques for developing such reduced-order models, in the context of a transitional channel flow.

The POD/Galerkin method has been used extensively for reduced-order modeling of fluid problems [13, 23, 31, 34, 35]. In this method, an empirical basis of orthonormal eigenfunctions is obtained from experimental or simulation data, and the Navier-Stokes equations are projected onto this basis. For fluid problems, this basis is optimal in terms of capturing the energy of the flow: the most significant modes are the ones that carry most of the kinetic energy. Although this method is applicable to many different types of flows, and computationally tractable for very large data sets, it can result in inaccurate low-order models since the most energetic modes are not always the most dynamically significant ones, as for example in the case of acoustic modes in cavity oscillations [31]. Improved POD models can in some cases be obtained by a careful selection of modes [35], removal of symmetry from the data using Fourier or traveling modes [14, 32] or inclusion of shift modes [25], but these techniques are often ad hoc and typically require extensive fine tuning. In addition, there is no standard method for incorporating actuation into POD models, and if using these models for control design, one needs to be careful to respect the region of validity of the model in the presence of actuation [36].

On the other hand, many tools from the control theory community have been developed for model reduction of linearized dynamics, often with a priori error bounds [26]. Here we focus on one of these methods, balanced truncation, introduced by Moore [24], and described in detail in standard references ([4, 26, 39]). Balanced truncation involves transforming a state-space system to a coordinate system where the states that respond most strongly to inputs (most controllable states) are also the states that have the most influence on future outputs (most observable states). In this *balanced realization*, weakly observable and controllable states can be truncated to form reduced-order models that capture well the input-output behavior of the system. The required coordinate transformation is known as the balancing transformation.

The main difficulty with using balanced truncation for large systems is computational expense. In a typical fluids problem, the number of states is on the order of 10^5 or more, and since computing balancing transformations involves solving Lyapunov equations for matrices of dimension $n \times n$ (where $n > 10^5$ is the number of states) as well as correspondingly large eigenvalue problems, traditional approaches to balanced truncations may not be feasible. Recently, snapshot-based methods for computing balanced truncation have been suggested, both to extend the concepts to nonlinear systems [18] and to large linear systems [30]. For systems with large numbers of outputs (e.g., if the output is the entire state), the snapshot-based method requires a large number of adjoint simulations, which still makes the problem computationally intractable. A procedure we call Balanced POD (BPOD) [30], uses an output projection to reduce the number of necessary adjoint simulations, and has been shown to be a computationally feasible approximation to balanced truncation for examples of two- and three-dimensional perturbations to laminar channel flow [14, 30].

Many recent works on transition in shear flows have focused on the large non-normal transient growth of exponentially stable linear perturbations to the laminar flow, which is thought to lead to the so-called 'subcritical' or 'bypass' transition [2, 3, 6, 16, 29, 33, 37]. A comprehensive treatment of the subject is given by Schmid and Henningson [33]. Suppression of this large growth is of interest in control applications for drag reduction. Control and estimation of linearized channel flow was studied by Högberg et al [12]. More recently, Åkervik et al [1] studied control of a cavity flow using global eigenmodes. In the works by Farrell and Ioannou [7] and Lee et al [20], balanced truncation was applied to linearized channel flow at particular wavenumbers where the standard algorithms are applicable, since the full system is one-dimensional.

The contribution of this work is the application of BPOD to transitional channel flow for

localized disturbances, without modeling on a wavenumber-by-wavenumber basis. Although a lot has been learned from studying 'canonical' classes of perturbations (streamwise vortices, oblique waves, Tollmien-Schlichting waves, simple localized disturbances [10]), the standard wavenumber-by-wavenumber analysis has its limitations when efficient implementation of closed-loop control is desired, since model reduction and control design would need to be performed at each wavenumber. The perturbations that arise in real flows often have complex three-dimensional structure with contributions at a wide range of wavenumbers as well as non-periodic or stochastic components. More importantly, for modeling of realistic actuation devices the wavenumber approach would be very complex except for the special case of actuators that excite only specific wavenumbers. It is therefore of interest to have a method where disturbances with complex structure are modeled without Fourier decomposition, in particular for controls applications where physically realizable localized actuators need to be considered. To the best of our knowledge, such low-order balanced truncation models have not been reported in literature. Although some success has been achieved for channel flow using control at all wavenumbers, the method we propose is able to extract dominant dynamics of the flow in physical space, resulting in simpler models. Moreover, this method for balanced truncation would be very advantageous for more complex geometries, such as spatially developing boundary layers.

We first obtain models of a single-wavenumber perturbation computed by Farrell [3] in order to validate our numerical methods, and then apply BPOD to a localized body-force actuator. The most desirable features of a reduced-order model are close approximation of the dynamics of the original system, inclusion of actuation and validity over a wide range of parameters, and we show that our BPOD models have those features. We also show that BPOD models capture the effects of actuation better than standard POD and that the considerable improvement in the capturing of the dynamics for BPOD is due to the non-orthogonal projection used.

The rest of this paper is organized as follows. In Section II, we briefly overview the two model reduction methods we use. In Section III, we describe the governing equations and our choice of inner product for the adjoint simulations of the system. In Sections IV and V, we present low-order models for a single-wavenumber optimal perturbation and a localized actuator and discuss their performance. Finally, in Section VI, we describe our conclusions and directions for future work.

II. MODEL REDUCTION VIA POD AND BALANCED POD

A. Proper Orthogonal Decomposition (POD)

We first give a brief overview of the POD/Galerkin method; details can be found in standard references [13, 34]. The idea of Galerkin projection is, given a system

$$\dot{x} = f(x), \qquad x(t) \in \mathcal{X}, \tag{1}$$

where $\mathcal{X}$ is a high-dimensional Hilbert space, to project onto a low-dimensional subspace $S \subset \mathcal{X}$. Proper Orthogonal Decomposition determines an orthogonal basis for such a subspace, which is obtained by solving the eigenvalue problem for XX^T where X is a matrix whose columns are simulation snapshots $x(t_k)$ at some times t_k. In this basis, we can represent the dynamics of $x(t)$ as

$$x(t) = \sum_{j=1}^{m} a_j(t)\theta_j, \tag{2}$$

where θ_j are the time-independent basis functions (POD modes) and $a_j(t)$ are the corresponding time coefficients, which are obtained from

$$\dot{a}_j = \langle \theta_j, f(r) \rangle . \tag{3}$$

A reduced-order model of order r can be obtained as a set of ODEs for the time evolution of these coeffecients by projecting the original system onto the most significant POD modes in the system (i.e. including only the first r modes where $r < m$). The method is applicable to both linear and nonlinear systems. For linear systems, the POD modes of data arising from the input-state impulse response are the most controllable modes of the linear system [30]. However, both controllability and observability are important for the input-output behavior of a system, and POD often fails to capture highly observable modes. On the other hand, balanced truncation does take into account both of these properties, and we describe this method next.

B. Balanced POD

Balanced truncation is a standard model reduction method [4, 24, 39] used for stable linear input-output systems of the form

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx \end{aligned} \tag{4}$$

where $u \in \mathcal{U} = \mathbb{R}^p$ is the vector of inputs, $y \in \mathcal{Y} = \mathbb{R}^q$ is the output, $x \in \mathcal{X} = \mathbb{R}^n$ is the state vector (although in general all three spaces can be complex as well), and A, B, and C are matrices of appropriate dimension. The idea of balancing is to find a change of coordinates in which the controllability and observability Gramians, defined by

$$W_c = \int_0^\infty e^{At} B B^+ e^{A^+ t}\, dt, \quad W_o = \int_0^\infty e^{A^+ t} C^+ C e^{At}\, dt, \tag{5}$$

are equal and diagonal. Here A^+, B^+ and C^+ define the corresponding adjoint system. It should be noted that in general $A^+ \neq A^T$, the two being equal only when the inner product used to derive the adjoint does not have an associated weight. It can be shown that balanced truncation does not depend on the choice of the inner product on the state space $\mathcal{X}$ (see Appendix A), although it does depend on the choices of inner products for $\mathcal{U}$ and $\mathcal{Y}$. One then truncates the least controllable and observable modes, corresponding to the smallest eigenvalues of these Gramians. A detailed description of the Balanced POD procedure, which is a computationally tractable procedure for finding such a transformation, is given in [30]. In this method, one begins by computing snapshots of the impulse-state response of the system in Eq. (4) and the adjoint system

$$\dot{z} = A^+ z + C^+ v \tag{6}$$

and stacking the direct and adjoint snapshots as columns of matrices X and Y (with appropriate quadrature weights [30]). One can show that the Gramians in Eq. (5) may then be approximated by *empirical Gramians* [18] $W_{c,e}$ and $W_{o,e}$, as

$$W_c \approx W_{c,e} = XX^+, \qquad W_o \approx W_{o,e} = YY^+. \tag{7}$$

The key idea in the method of snapshots is to compute the transformation that balances the empirical Gramians (or at least the dominant directions of this transformation) without actually computing the Gramians themselves, whose dimension is large. In this respect, the method of snapshots for BPOD [30] resembles the method of snapshots introduced by Sirovich [34] for more efficient computation of POD modes. To compute the balancing transformation, one computes the singular value decomposition (SVD) of the matrix Y^+X (see Appendix A for a discussion of Y^+):

$$Y^+ X = U \Sigma V^T, \tag{8}$$

from which the balancing transformation Φ and its inverse Ψ are found by

$$\Phi = XV\Sigma^{-1/2}, \quad \Psi = YU\Sigma^{-1/2}. \tag{9}$$

The columns of Φ are the *balancing modes* and the columns of Ψ are the *adjoint modes*, and the two sets of modes are biorthogonal. The entries of the diagonal matrix Σ are known as the Hankel singular values (HSVs). The non-orthogonal projection onto the basis of balancing modes using the adjoint modes is also known as Petrov-Galerkin projection.

Note that a different procedure for approximating balancing transformations has also been used in [38], in which the Gramians are separately reduced (that is, low-rank approximations of $W_{c,e}$ and $W_{o,e}$ are first constructed, and then the balancing transformation for the rank-reduced Gramians is computed by an unspecified algorithm). However, this procedure is more computationally intensive than our procedure, and also gives worse results [30], since almost-uncontrollable modes may be strongly observable, so should not be truncated.

C. Output projection

If the number of outputs of the system is large, as in a typical fluids problem (where $n = q$, i.e. the output is the full state), the computation of the adjoint simulations of the system given by Eq. (6) may not be tractable, since one simulation is needed for each component of the output. A way to reduce the number of system outputs is to first project the output onto a low-dimensional subspace, i.e., taking $\tilde{y} = P_s C x$, where P_s is an orthogonal projection onto a s-dimensional subspace of $\mathcal{Y}$, as suggested in [30]. The system is now of the form:

$$\begin{aligned} \dot{x} &= Ax + Bu \\ \tilde{y} &= P_s C x \end{aligned} \tag{10}$$

where s is the rank of the output projection. The projection P_s that minimizes the 2-norm of the difference between the original transfer function and the output-projected transfer function is given simply by the POD of the set of impulse-state responses [30]. This projection can be written as $P_s = \Theta_s \Theta_s^+$, where columns of $\Theta_s : \mathbb{R}^s \to \mathcal{Y}$ are POD modes. Another way to write the system is as follows:

$$\begin{aligned} \dot{x} &= Ax + Bu \\ \hat{y} &= \Theta_s^+ C x \end{aligned} \tag{11}$$

Here, the outputs of the system are just the coefficients of the POD modes of the system impulse response and $\hat{y} \in \mathbb{R}^s$. This s-dimensional output carries the same information as the n-dimensional output $\tilde{y}$, which can be shown by Parseval's theorem. The corresponding adjoint system can now

be written as

$$\begin{aligned} \dot{z} &= A^{+}z + (\Theta_s^{+}C)^{+}v \\ w &= B^{+}z. \end{aligned} \tag{12}$$

Note that if the output is the full state, so that $C = I$, and the adjoint is defined with respect to the standard L_2 inner product, the initial conditions of the adjoint simulations are just the POD modes (columns of Θ_s). In practical computations, depending on the choice of inner product used in defining the adjoint system, and on the numerical quadrature method (for example, if the computations are done using Chebyshev polynomials) the matrix Θ_s^{+} is usually just the matrix Θ_s pre-multiplied by a matrix of inner product weights.

The idea of Balanced POD is to compute the snapshot-based balanced truncation of the system in Eq. (11) instead of Eq. (4), so that only s adjoint simulations are needed. It is easily shown that the systems in Eqs. (10) and (11) have the same observability Gramian using the fact that for any projection P, we have $P^2 = P$. Transforming Eq. (11) to balanced coordinates and writing $x = \Phi_1 a$ is obtained as follows:

$$\begin{aligned} \dot{a} &= \Psi_1^{+}A\Phi_1 a + \Psi_1^{+}Bu \\ y &= \Theta_s^{+}C\Phi_1 a. \end{aligned} \tag{13}$$

The inverse transformation matrix $\Psi_1 : \mathbb{R}^r \to \mathcal{X}$ and the transformation matrix $\Phi_1 : \mathbb{R}^r \to \mathcal{X}$ are $n \times r$, r being the number of states we want to retain in the system, which we will refer to as the rank of the model. Note that $r \leq p$, where p is the number of non-zero HSVs. For simplicity, we will assume from now on that $C = I_n$, i.e. the output of the original system is the full state (this is the case in fluid simulations in which we need to know the entire flow field). We can then represent the output of Eq. (13) as $y = \Theta_s^{+}\Phi_1 z$, which is now the vector of time coefficients of the s standard POD modes obtained from the impulse response of the system. For fluids systems the full flow field output of the model can be recovered from these coefficients and the corresponding modes. For a given dimension of the output projection, all BPOD models will have s outputs regardless of the model rank r, while the number of POD model outputs is equal to r at each rank. The effect of output projection on model performance will be illustrated in Sec. IV and Sec. V.

III. APPLICATION TO TRANSITIONAL CHANNEL FLOW

A. Governing equations

For shear flows, the linearized equations may be conveniently written in terms of the wall-normal velocity v and the wall-normal vorticity η (see, for instance, [33]). The other variables (e.g., streamwise and spanwise velocities u and w) may then be computed using the continuity equation $\partial_x u + \partial_y v + \partial_z w = 0$ and the definition of wall-normal vorticity. In these coordinates, the linearized (nondimensional) equations have the form

$$\left[(\partial_t + U\partial_x)\Delta - U''\partial_x - \frac{1}{Re}\Delta^2\right] v \;=\; 0 \tag{14}$$

$$\left[\partial_t + U\partial_x - \frac{1}{Re}\Delta\right] \eta \;=\; -U'\frac{\partial v}{\partial z}. \tag{15}$$

Here, $Re = U_c\delta/\nu$ is the Reynolds number, where ν is the kinematic viscosity, δ is the half-width of the channel, and $\Delta = \partial_x^2 + \partial_y^2 + \partial_z^2$ is the Laplacian. U_c is a characteristic velocity, which for linearized channel flow is the centerline velocity of the laminar profile $U(y)$. The prime indicates differentiation with respect to y. The first equation is the Orr-Sommerfeld equation and the second one is known as the Squire equation. It was first shown numerically by Orszag [27] that the Orr-Sommerfeld equation for channel flow is stable up to $Re \approx 5772$, when an exponentially unstable eigenmode first arises. The Squire equation has stable eigenmodes for all values of Re. Still, complex behavior due to the non-normality exists for stable eigenmodes. The term on the right hand side of the Squire equation represents tilting of the spanwise component of the vorticity of the mean flow (which here is just U') by the strain rate $\partial v/\partial z$ [3], which gives rise to wall-normal vorticity. In the limit of high Reynolds number, the perturbation growth is dominated by this process, in particular for streamwise-constant perturbations. While the system also exhibits phenomena such as degeneracies and resonances [8, 11], non-normality has been shown to have a dominating effect on the energy growth [28].

In operator form, we can represent the equations using more compact notation as follows:

$$\frac{\partial}{\partial t}\begin{bmatrix} -\Delta & 0 \\ 0 & I \end{bmatrix}\begin{bmatrix} v \\ \eta \end{bmatrix} = \begin{bmatrix} L_{OS} & 0 \\ -U'\partial_z & L_{SQ} \end{bmatrix}\begin{bmatrix} v \\ \eta \end{bmatrix} \tag{16}$$

where

$$L_{OS} = U\partial_x\Delta - U''\partial_x - \frac{1}{Re}\Delta^2$$
$$L_{SQ} = -U\partial_x + \frac{1}{Re}\Delta$$

are the Orr-Sommerfeld and Squire operators, respectively. If we define

$$A = \begin{bmatrix} -\Delta & 0 \\ 0 & I \end{bmatrix}^{-1} \begin{bmatrix} L_{OS} & 0 \\ -U'\partial_z & L_{SQ} \end{bmatrix} \tag{17}$$

with no-slip boundary conditions, we can write the system in standard state-space form:

$$\begin{aligned} \dot{x} &= Ax + Bu_1 + Fu_2 \\ y &= Cx \end{aligned} \tag{18}$$

where B and F represent the spatial distributions of the actuators and disturbances respectively, with $u_1(t)$ and $u_2(t)$ being the corresponding input vectors (the time-dependent amplitudes of the columns of B and F). The actuation and the disturbances are equivalent mathematically as they are both inputs to the system. We note here that the impulse-state responses are given by $x_1(t) = e^{At}B$ and $x_2(t) = e^{At}F$, and the adjoint system impulse-state responses for the full system are given by $z(t) = e^{A^+t}C^+$. Therefore, to obtain the POD basis needed for BPOD, we simulate the system given by Eq. (14) with a given perturbation or actuator as initial condition until the response has decayed to negligible levels, so that the matrix XX^T that can be formed from the snapshots will closely approximate the controllability Gramian given by Eq. (5), where the integral extends to infinite time. Of course, computation of the matrix XX^T is intractable for very large systems, so we compute POD via the method of snapshots, forming the smaller matrix X^TX and following the procedure described in Sec. II A.

B. Inner product on the state space

To determine the corresponding adjoint equations, one first needs to define an inner product on the vector space $\mathcal{X}$ of flow variables (v, η). Since balanced truncation is independent of the choice of inner product used to define the adjoint (see Appendix A), we may choose an inner product which is convenient for numerical computations. Let us define the inner product

$$\langle (v_1, \eta_1), (v_2, \eta_2) \rangle_M = \int_\Omega (-v_1\Delta v_2 + \eta_1\eta_2)\, dx\, dy\, dz, \tag{19}$$

where Ω denotes the fluid volume. Note that, letting $M : \mathcal{X} \to \mathcal{X}$ denote the matrix operator on the left hand side of Eq. (16), this is just the L_2 inner product of (v_1, η_1) with $M(v_2, \eta_2)$. This inner product is different from the standard energy inner product used in analyzing perturbations through Fourier decomposition [3, 8], as there is no re-scaling at each wavenumber.

With this definition of the inner product, the adjoint equations are easily found by integration by parts:

$$\frac{\partial}{\partial t}\begin{bmatrix} -\Delta & 0 \\ 0 & I \end{bmatrix}\begin{bmatrix} v \\ \eta \end{bmatrix} = \begin{bmatrix} L_{OS}^* & U'\partial_z \\ 0 & L_{SQ}^* \end{bmatrix}\begin{bmatrix} v \\ \eta \end{bmatrix} \tag{20}$$

where

$$L_{OS}^* = -U\partial_x\Delta - 2U'\partial_x\partial_y - \frac{1}{Re}\Delta^2$$
$$L_{SQ}^* = U\partial_x + \frac{1}{Re}\Delta.$$

C. Inner product on the output space

Although the time evolution of the linearized disturbances is fully described by the wall-normal velocity-vorticity formulation, the output of the system can be chosen to be in different variables. When using POD, the choice of inner product can have a large impact on the results. If the output of our system is only the velocity-vorticity field, the standard L_2 inner product can be used. For our system, since the other two velocity components can easily be recovered using continuity and the definition of vorticity, we can choose the full velocity field to be the output, and use the energy inner product given by

$$\langle \mathbf{u_1}, \mathbf{u_2} \rangle = \int_\Omega (u_1 u_2 + v_1 v_2 + w_1 w_2)\, dx\, dy\, dz, \tag{21}$$

This choice is more intuitively appealing, since the POD modes for the output projection will capture the true kinetic energy of the perturbation. We therefore define the output space $\mathcal{Y}$ in our system as the space $\mathbb{R}^n$ together with the inner product defined by Eq. (21). We note here that the space $\mathcal{X}$ is also $\mathbb{R}^n$, though endowed with a different inner product (the M-inner product described in the previous section).

D. Numerical Methods

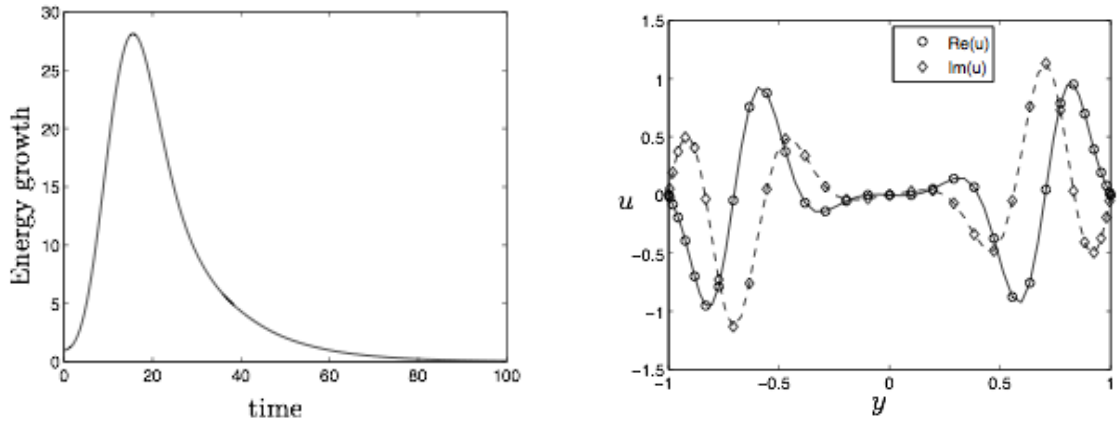
The simulations were performed using a linearized version of a fully nonlinear DNS code using the spectral method described by Kim et al. [17], with periodic boundary conditions in the streamwise and spanwise directions. The linearized code was verified against the analytic time evolution of Orr-Sommerfeld eigenfunctions and optimal perturbations, and the resolution of each simulation was checked by varying the time step and grid resolution. The size of the computational box was $2\pi \times 2\pi$ in the streamwise and spanwise directions for all simulations. Standard LAPACK routines were used for the computation of POD and balanced POD modes, as well as for the comparison of subspaces. The reduced-order models were integrated using the standard fourth-order Runge-Kutta scheme. All computations were done using Fortran 90 and MATLAB. A code by S.C.Reddy [33] was used to compute the initial conditions for the single-wavenumber perturbations. The integration weights derived by Hanifi et al [9] were used for the computation of inner products on the Chebyshev grid.

IV. RESULTS: SINGLE-WAVENUMBER PERTURBATIONS

We start by investigating the system given by Eq. (18) without actuation and only in the presence of disturbances (without the Bu_1 term). In order to validate the numerical methods, we first obtain BPOD models from three-dimensional simulations of simple and well-known single-wavenumber perturbation cases, described by Butler and Farrell [3] and also investigated by Schmid and Henningson [33]. The general form of such disturbances is given by

$$q(x, y, z, t) = \hat{q}(y, t)e^{(i\alpha x + i\beta z)}, \tag{22}$$

where $\hat{q}(y, t) = [v(y, t)\ \eta(y, t)]^T$. The standard approach to such perturbations is to compute the time evolution of $\hat{q}(y, t)$, which fully describes the system, since the velocity components u and w can easily be computed. For this one-dimensional problem, standard algorithms for computing balanced truncation are computationally tractable. Therefore, we are able to compare the models resulting from exact balanced truncation (which for the 1-D case can be computed in MATLAB using standard algorithms [19]) to BPOD models obtained from three-dimensional simulations of the real part of the full field, $Re\{q(x, y, z, t)\}$ at a particular wavenumber pair (α, β) on a large grid, similar to the comparison made by Rowley [30] for a streamwise-constant perturbation. We note that for a given wavenumber pair the comparison between BPOD and exact balanced

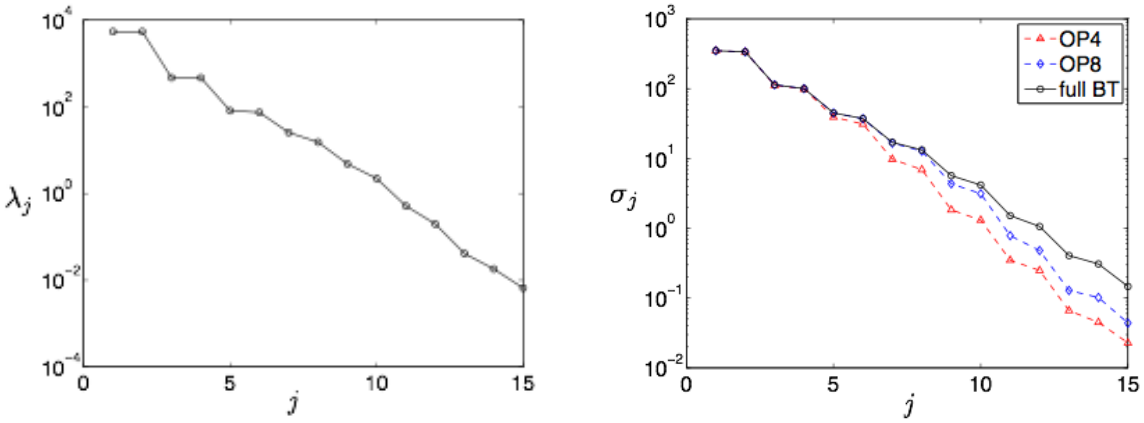


(a) (b)

FIG. 1: (a) Kinetic energy growth for the optimal perturbation at wavenumber $\alpha = 1, \beta = 1$ at $Re = 1000$. (b) The $\alpha = 1, \beta = 1$ optimal perturbation, showing streamwise velocity u (complex).

truncation can be done only using 1-D simulations, but we also performed 3-D simulations in order to verify our codes. We also note that, since the outputs of the output-projected system and the reduced-order models are coefficients of POD modes, the C matrix in (18) was modified so that the output of the full system is in the POD basis as well. This way, a meaningful comparison between the balanced truncation of the full system and BPOD is obtained.

The initial conditions were computed using the method described by Reddy and Henningson [28] and their energy growth was verified against values reported in that work. While streamwise-constant perturbations exhibit the largest energy growth, three-dimensional perturbations exhibit more interesting dynamics. (Here, by "three-dimensional," we mean that the perturbations have components in both streamwise and spanwise directions, although the problem can still be treated as 1-D in the wall-normal direction as described above.) We focus on the $\alpha = 1, \beta = 1$ perturbation at $Re = 1000$, whose energy growth is shown in Fig. 1. The computational grid used in the three-dimensional simulation was $16 \times 65 \times 16$, corresponding to 33280 states in the system given by Eq. (18). Balanced truncation of the one-dimensional problem with 65 Chebyshev modes is easily and accurately computed using the algorithm described in [30] so that BPOD performed on the large system can be compared to exact balanced truncation.



(a) (b)

FIG. 2: (a) The first 15 POD eigenvalues for $\alpha = 1, \beta = 1$ initial perturbation at $Re = 1000$. (b) The first 15 Hankel singular values (HSVs) for: four-mode ($\triangle$) and eight-mode ($\diamond$) output projections and full balanced truncation ($\circ$) for the same case.

A. Mode subspaces

It was found that 500 equally spaced snapshots are sufficient for accurate computation of the POD modes, since for a larger number of snapshots with finer spacing there is no considerable change in the eigenvalue spectrum or the corresponding modes. We see from Fig. 2 that the most significant eigenvalues and the corresponding modes typically come in pairs, representing traveling structures that are 90 degrees out of phase. The first pair of modes contains 90.45% of the energy, while the first three pairs contain 99.6% of the energy. For the balanced POD models, a four-mode and eight-mode output projections were chosen, corresponding to respectively 98.3% and 99.9% of total energy contained in the POD modes.

We also notice that the HSVs (Fig. 2) come in pairs, indicating that the most significant modes in the BPOD mode basis are again traveling structures similar to the standard POD modes. It is important to include these pairs of modes in the reduced-order models, as stability of the models for balanced truncation is guaranteed only if $\sigma_r > \sigma_{r+1}$ where r is the rank of the model [4]. While for standard POD modes there is no such requirement, mode pairs should always be included in the models on physical grounds. We also notice that the number of HSVs for each output projection that is equal to the full balanced truncation HSVs is approximately equal to the output projection rank. The same observation was made by Rowley [30], although there is no proof of this property

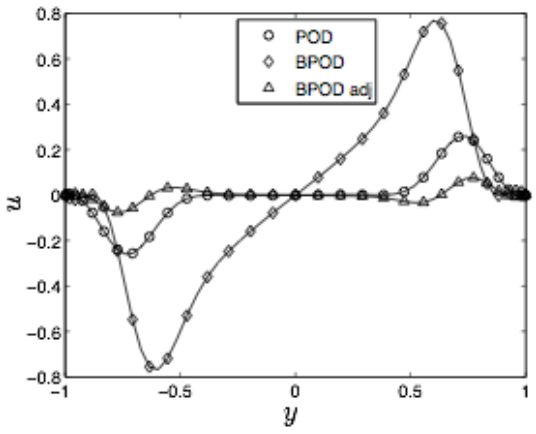
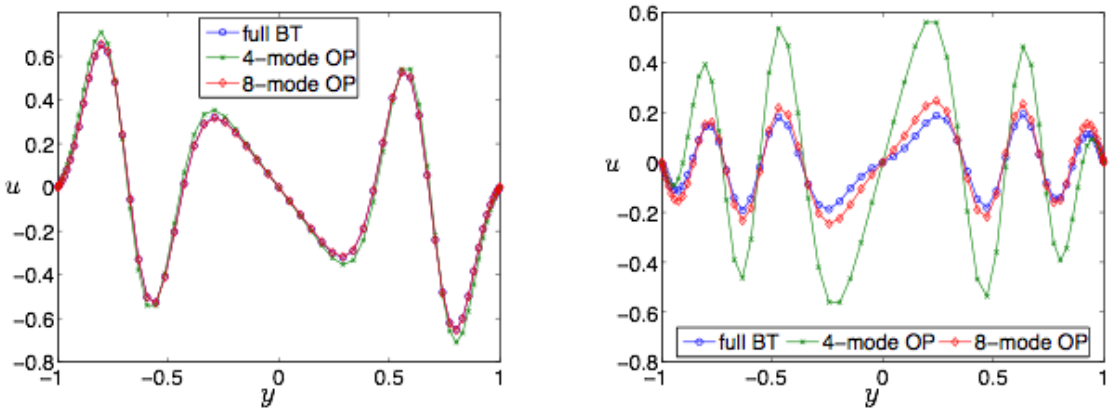


FIG. 3: Streamwise velocity for the first POD mode, balancing mode and adjoint mode for the $\alpha = 1, \beta = 1$, $Re = 1000$ initial condition.

at this point.

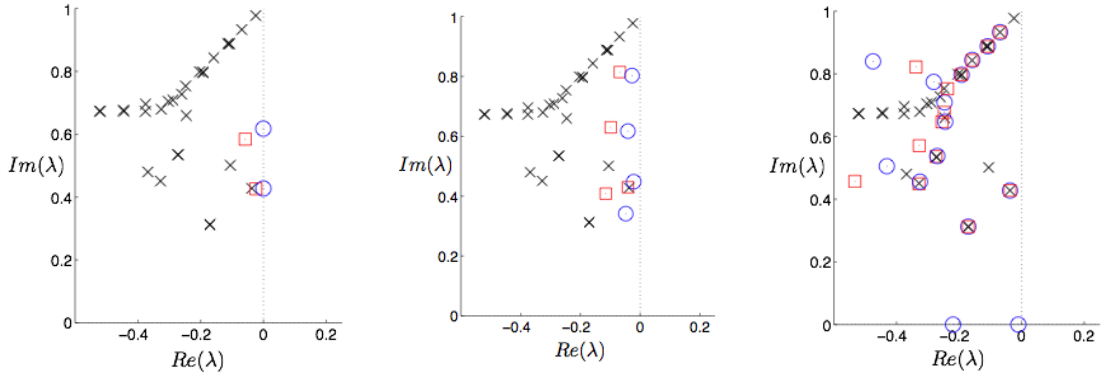
The first POD mode is shown in Fig. 3 together with the first balancing and adjoint modes from a four-mode output projection. Figure 4 shows the streamwise velocity of the sixth and tenth balancing modes, illustrating the effect of the choice of output projection rank. The first four balancing modes from BPOD look identical for both output projections, while the sixth mode is not very accurately captured by a four-mode output projection. Both output projections do not capture very accurately the higher modes such as the tenth mode, which is also illustrated by the HSVs in Fig. 2. As we show below, this inaccuracy does not significantly affect model performance, since these higher modes are not very significant dynamically.

In this single-wavenumber case, the exact eigenvalue spectrum of the A matrix from Eq. (17) at a given Reynolds number can easily be computed. We note here that the eigenvalues of the matrix A and therefore the poles of the corresponding transfer function are independent of the initial condition (which is just the B matrix for our impulse response simulations). Figure 5 shows the spectra of the full operator and three reduced-order models of different rank for the $\alpha = 1, \beta = 1$ perturbation. Since the spectra are symmetric about the real axis, we only show the upper half of the complex plane. We see that, while the representation of the full spectrum improves for both methods as the rank increases, BPOD captures more accurately some of the most slowly decaying eigenvalues, which have the most influence on the dynamics of the system. For the rank four model, the standard POD model appears to be marginally stable, while the BPOD model approximates closely the eigenvalue closest to the origin. At higher order, the standard POD



(a) (b)

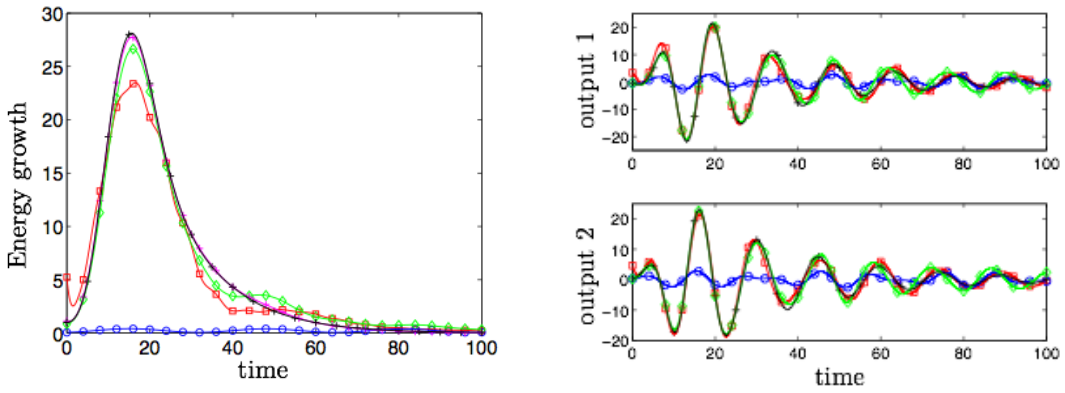
FIG. 4: Streamwise velocity for (a) the sixth balancing mode and (b) the tenth balancing mode for BPOD with two different output projections and for full balanced truncation.



(a) (b) (c)

FIG. 5: Spectrum of the full operator and reduced-order models for rank (a) 4, (b) 8, (c) 30. The BPOD modes are from the eight-mode output projection. Symbols: BPOD (□), POD (◯), full operator (+) Only the most important part of the full spectrum is shown.

models improve and capture approximately the same eigenvalues as the BPOD models of the same rank. It is also important to notice that some of the eigenvalues of the full system are never captured by the reduced-order models. These eigenvalues correspond to uncontrollable eigenmodes of the full system, and can never be excited by this particular perturbation.



(a) (b)

FIG. 6: (a) $\alpha = 1, \beta = 1$ optimal perturbation at $Re = 1000$, eight-mode output projection, 4-mode and 8-mode models. Full simulation (+), 4-mode POD (◯), 4-mode BPOD (□), 8-mode POD (◇), 8-mode BPOD (+) (b) First two outputs, symbols as defined in (a).

B. Impulse response

We next compare the impulse response of the system to that of the reduced-order models. The impulse response of a linear system is important, since the response of the system to any input can be found from the convolution of the impulse response with the input. Figure 6 shows the capturing of the growth of kinetic energy by POD and BPOD models, as well as the first two outputs of the reduced-order models. The poor performance of POD at low orders for the traveling structure perturbation is evident. Even the eight-mode POD model, which captures the energy growth well, does not accurately capture the phase of the oscillations.

Figure 7 shows the 2-norms of the error between the impulse responses of the reduced-order models and the full simulation $\|G - G_r\|$, normalized by the 2-norm of the impulse response of the full simulation. This figure is a clear demonstration of the effect of output projection. A four-mode output projection means that we are effectively performing balanced truncation on the dynamics of the first four POD modes of the full system. The dashed lines in the figure indicate the 2-norms of the error between the full dynamics of the output-projected system and the full system. As the rank of the BPOD models is increased, the dashed lines, which are the limit of accuracy, are reached fast. As already seen in Fig. 6, for very low-order models, standard POD is clearly outperformed by balanced POD. We see that ten POD modes are needed to match the performance of a four-mode output projection BPOD model at the same rank. The slow improvement in POD

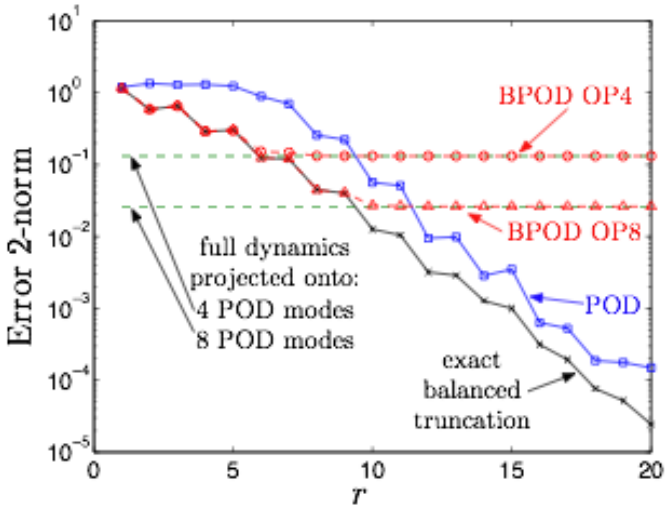
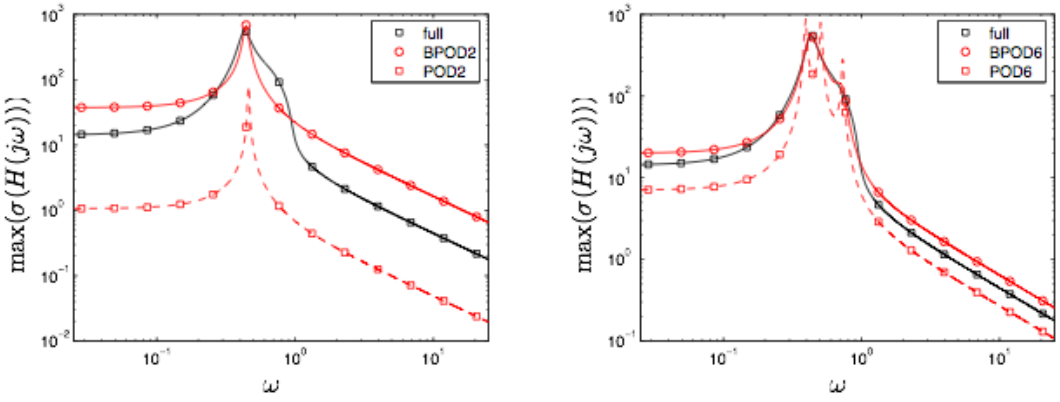


FIG. 7: Error 2-norms for the $\alpha = 1, \beta = 1, Re = 1000$ perturbation for full balanced truncation, POD and BPOD at two output projections.

model performance indicates that the dynamics of the perturbation can not be represented only by retaining the first few most controllable (POD) modes. Adding new BPOD modes beyond rank eight and ten (for four- and eight-mode output projections, respectively) does not improve the model performance noticeably, since the dynamics of the output-projected system is already captured fully. It is also important to note that the performance of the BPOD models is identical to that of full balanced truncation almost until the rank at which BPOD model error norms level off due to the output projection. This indicates that the higher balancing modes which are not computed accurately due to the approximation inherent in the output projection (such as those shown in Fig. 4) do not significantly influence the reduced-order model performance, the main limitation being the capturing of the full system by the output projection.

C. Frequency Response

The frequency response encompasses system behavior over the complete range of possible forcing, and is perhaps the best indication of overall system performance. Therefore, from the control designer's point of view, having a low-order model that represents well the frequency response of the original system is of key importance. Frequency response of single-wavenumber perturbations was investigated by Schmid and Henningson [33] using the resolvent norm, where



(a) (b)

FIG. 8: Frequency response of the models for the $\alpha = 1, \beta = 1$ perturbation. (a) 2-mode POD (dashed), 2-mode BPOD (full), (b) 6-mode POD (dashed), 6-mode BPOD (full).

at each frequency the maximum amplification over all initial conditions is computed. Here the frequency response of the system with a given actuator or perturbation is of interest.

A standard way of representing synthesized frequency response for MIMO (multiple input multiple output) systems is a plot of the maximum singular value of the transfer function matrix $\max(\sigma(H(i\omega)))$ as a function of frequency, also known as a singular value Bode plot. Fig. 8 (a) and (b) shows such plots for the $\alpha = 1, \beta = 1$ perturbation and is a clear demonstration of the advantages of BPOD for capturing the dynamics of the system. We see that even for a two-dimensional model the resonant peak is captured well by the model, while for standard POD the peak is very narrow, with very low response at other frequencies. This behavior is typical of balanced truncation, as shown in [4] - the first modes to be captured are the ones which are most significant dynamically, while the correct response is gradually built up in less significant frequency bands as more modes are added. For standard POD models, on the contrary, the response improves incrementally at all frequencies as more modes are added and a higher number of modes is needed to accurately capture the resonant peak. For ten mode models, both standard POD and BPOD capture well the frequency response (not shown in the figure), with BPOD frequency response being almost indistinguishable from the full system one. Standard POD frequency response also includes spurious non-physical peaks at low order, which correspond to eigenvalues very close to the imaginary axis for low order of truncation, as seen in Fig. 5 (a).

We are also interested in the worst-case error between the reduced-order model and the full

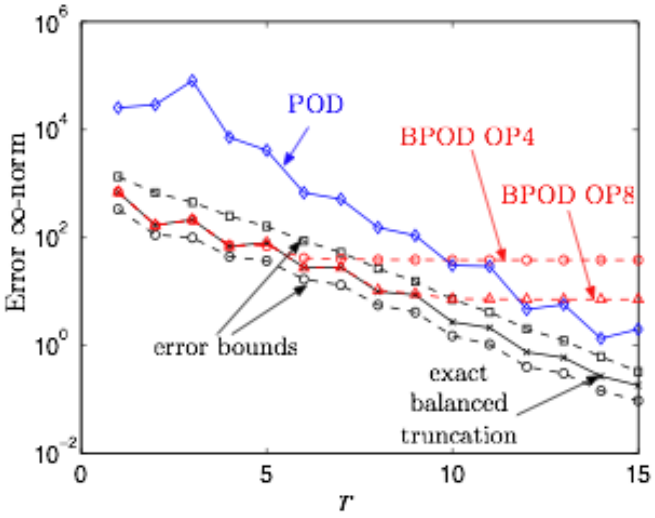


FIG. 9: Infinity error norms for POD, exact balanced truncation and BPOD with the infinity error bounds.

simulation, which is known as the infinity norm of the system. Balanced truncation has apriori error bounds for the infinity norm [4]. The error H_∞ lower bound for any reduced-order system is

$$\|G - G_r\|_\infty \geq \sigma_{r+1}, \tag{23}$$

where $G(s)$ is the full model, $G_r(s)$ is the reduced-order model with state dimension r and σ_j is the j-th Hankel singular value (in decreasing order). The upper bound for the error is given by

$$\|G - G_r\|_\infty \leq 2\Sigma_{j=r+1}^{n}\sigma_j. \tag{24}$$

The upper bound on the error can be very close to the lower bound if the HSVs decrease fast. Figure Fig. 9 shows the infinity norm of the error transfer function between the full system and the reduced-order model as a function of model rank for the first fifteen orders of truncation. The infinity norms for exact balanced truncation lie within the theoretical bounds given by Eqs. (23) and (24), while for BPOD, for each of the two output projections, the norms stay within bounds up to approximately the rank of the output projection, analogous to the two-norms in Fig. 7. The infinity norms for standard POD at low rank are considerably higher than those for balanced truncation and BPOD, corresponding to the frequency responses shown in Fig. 8 (a) and (b).

D. Variation of Reynolds number

Another very desirable feature of a reduced-order model is good performance for off-design values of the system parameters. We would like the models to remain valid for a wide range of the model parameters, or at least for the range appropriate for the physical application of the model. The only parameter we are considering in our models is the Reynolds number, so the response of models was compared to the full simulation when Re is changed. Separating the operators from Eq. (16) into convective and diffusive parts, we can re-write the state-space equation as

$$\dot{x} = A_{\text{conv}}x + \frac{1}{Re}A_{\text{diff}}x + Bu. \tag{25}$$

We can then separately project the matrices A_{conv} and A_{diff} as in Eq. (13) at any Reynolds number onto the POD and BPOD modes obtained at $Re = 1000$, the B matrix being just the initial condition at $Re = 1000$. Figure 10 shows the performance of 12-mode standard POD and BPOD models when the value of Re in Eq. (25) was changed to 2000 and the impulse response of the resulting models was compared to the impulse response of the full system. This rank of the model was chosen since both models perform well at the design condition of $Re = 1000$. We see that the BPOD model eigenvalues stay closer to the full simulation eigenvalues (which move as well), and also remain in the left half of the complex plane, while for $Re = 2000$ the standard POD model becomes unstable. This indicates a greater range of validity for BPOD models and better stability at off-design conditions than standard POD.

V. RESULTS: THREE-DIMENSIONAL LOCALIZED ACTUATOR

We next consider a localized body force actuator in the center of the channel which cannot be described by a one-dimensional problem. This case corresponds to Eq. (18) without the Fu_2 term, with the input matrix B representing the velocity field in Fig. 11. Individual localized disturbances to channel flow were investigated by Henningson et al [10]. Since balanced truncation involves the approximation of the system's Gramians (although in BPOD we do not actually compute the Gramians themselves), we are interested in following both the forward and adjoint impulse-state responses until all transients have completely decayed. The computational box necessary for following individual localized disturbances long enough in time would be very large, and we instead consider a periodic array of small disturbances in the channel center. It should be noted that the

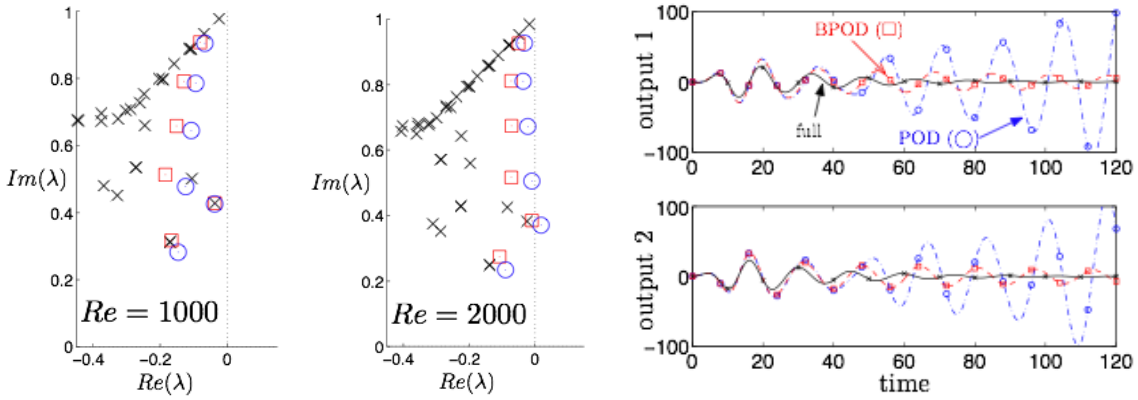


FIG. 10: Left: The comparison of spectra of the full operator at $\alpha = 1, \beta = 1$ to the spectra of rank 12 models as the value of Re is increased to 2000. Right: The performance of corresponding reduced-order models at $Re = 2000$ - first two outputs. See text for detailed description. Symbols: full ($\times$), BPOD ($\square$), POD ($\bigcirc$).

behavior of this periodic array can be quite different from the behavior of a single localized disturbance, in particular considering the energy growth, since the periodic array quickly develops into a streamwise-constant vortex. The exact form of the initial condition considered here is

$$v(x, y, z, 0) = A\left(1 - \frac{r^2}{\alpha^2}\right) e^{(-r^2/\alpha^2 - y^2/\alpha_y^2)}(\cos(\pi y) + 1) \tag{26}$$

where $x_c, 0, z_c$ are the coordinates of the center of the computational domain and $r^2 = (x - x_c)^2 + (z - z_c)^2$. The wall-normal vorticity is zero. This form was picked in order to satisfy the condition that the mean perturbation velocity is zero at each wavenumber. The $(\cos(\pi y)+1)$ term was added to make the field satisfy exactly the boundary conditions $v(\pm 1) = v_y(\pm 1) = 0$. The amplitude A was set to 1 for this simulation, and the parameters α and α_y were set to $\alpha = 0.7$ and $\alpha_y = 0.6$. The Reynolds number chosen for this simulation was $Re = 2000$. The traveling structure rapidly develops into a streamwise-constant structure, since the growth of wall-normal vorticity results in the development of streamwise streaks (see Fig. 11).

The grid size was $32 \times 65 \times 32$, corresponding to 133,120 states for the full (v, η) system. The simulation was ran for 1200 dimensionless time units ($t = t^d U_c/\delta$), and the timestep used was $\Delta t = 0.004$. During this time, the energy of the initial disturbance decayed to approximately 1.5 percent of its initial value. The POD modes were taken over 1000 snapshots, with fine spacing between snapshots for the initial period in order to capture the traveling structures well and larger spacing once the streamwise structure was developed, after it was verified that POD eigenvalues

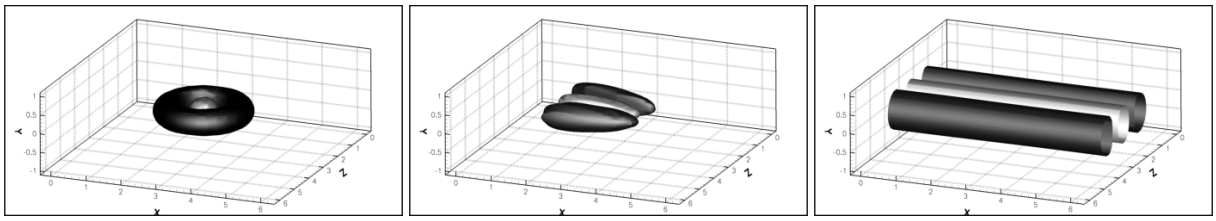
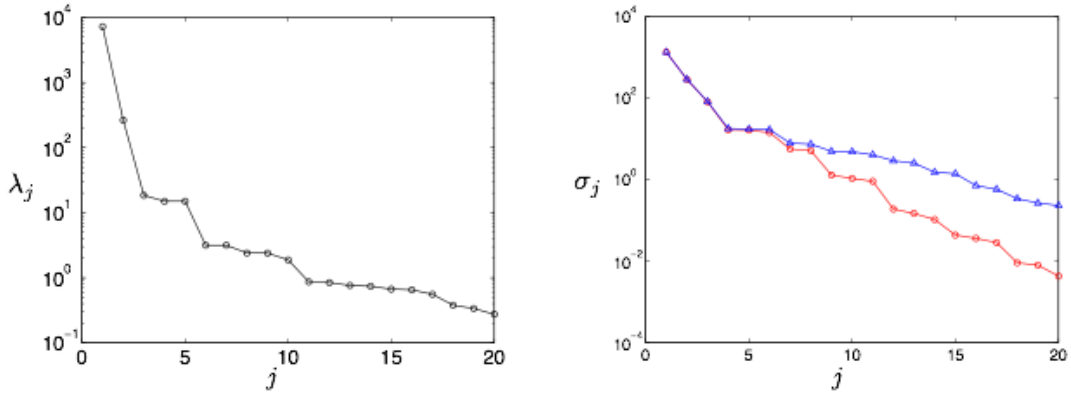


FIG. 11: The development of the wall-normal velocity of the perturbation given by Eq. (26) at $t = 0$ (left), $t = 14$ (middle) and $t = 160$ (right) which corresponds to the maximum energy growth. Positive velocity is light, and negative velocity is dark.



(a) (b)

FIG. 12: (a) The first 20 POD eigenvalues for the Gaussian-like disturbance impulse response. (b) The first 20 HSVs for five-mode (○) and ten-mode (△) output projections.

and the corresponding modes do not change significantly if more snapshots are used. Fig. 12 (a) shows the POD eigenvalues of the impulse response. The first five modes contain 99.72% of the perturbation energy, and the first ten modes contain 99.9% of the energy. In this case the spectrum contains both streamwise-constant (and nearly-streamwise constant modes) as well as traveling structures due to the initial transient. The first three modes are streamwise-constant structures, while the fourth and the fifth modes correspond to a traveling structure, which accounts for only 0.40% of the total energy. Modes one, four and five are shown in Fig. 13.

Next, the adjoint simulations were ran and the BPOD procedure was performed on a five-mode output projection, containing only the most important traveling structure, as well as on a ten-mode output projection. These ranks were chosen due to large drops in energy significance after the fifth and tenth mode, as shown in Fig. 12 (a). Fig. 12 (b) shows the HSVs for these two output

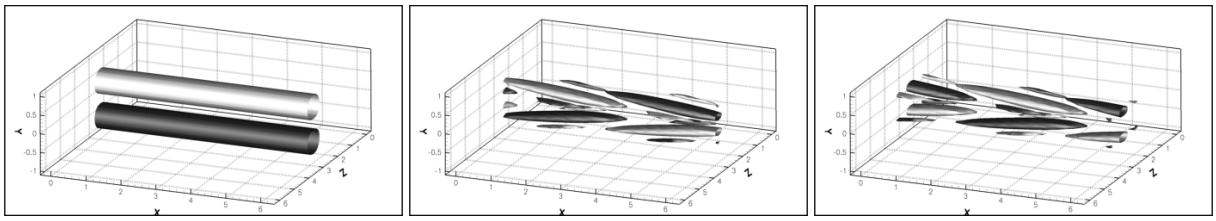


FIG. 13: The first, fourth and fifth POD modes for the localized actuator, showing streamwise velocity. The iso-surfaces show half of the maximum (light) and minimum (dark) value.

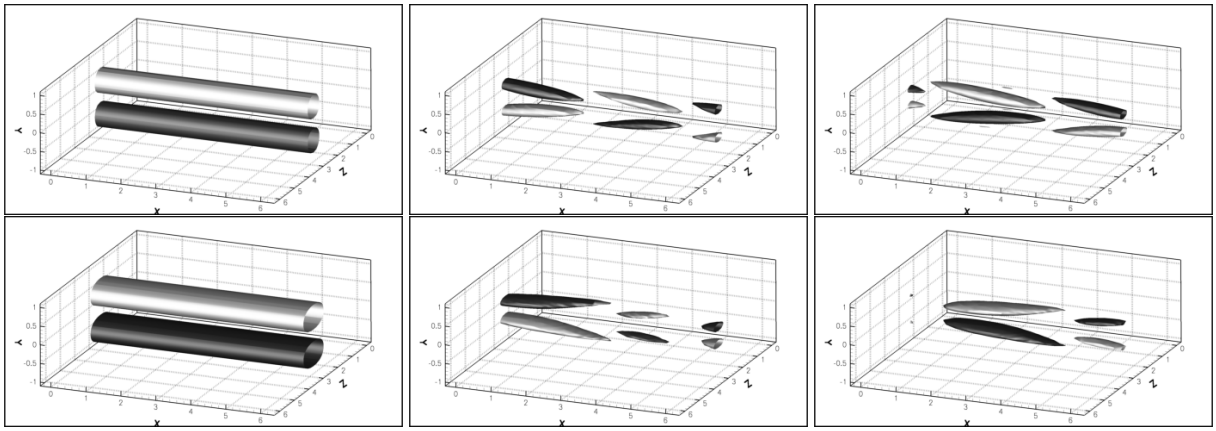


FIG. 14: Top row: primal modes one, four and five from balanced POD, showing streamwise velocity for the localized perturbation. The modes are from a five-mode output projection. The iso-surfaces show half of the maximum (light) and minimum (dark) value. Bottom row: the corresponding adjoint modes. Note the similarity between the primal modes and the corresponding POD modes in Fig. 13.

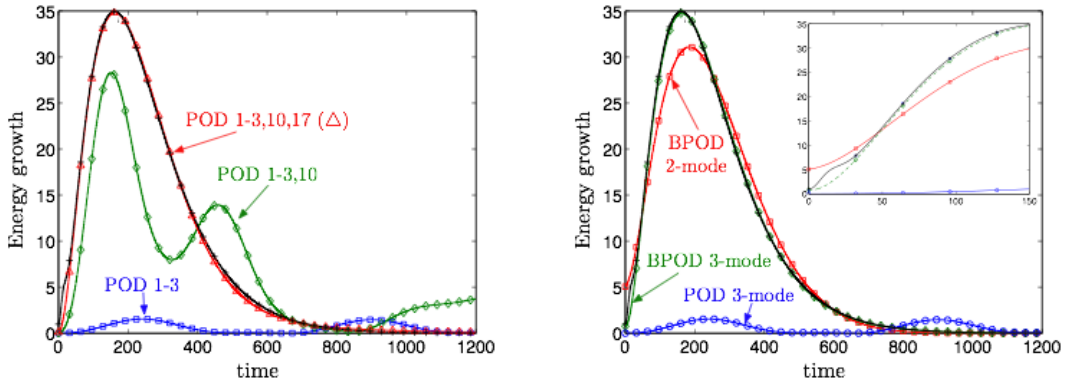
projections. We notice that the HSVs are equal for the pairs of modes 4–5 and 7–8 for five-mode output projection, corresponding to traveling structures in the basis of BPOD modes. Even more interestingly, for the ten-mode output projection, HSVs for the modes 4–6 are equal. Although the stability of balanced truncation models is guaranteed only when $\sigma_{r+1} < \sigma_r$, where r is the number of states retained [4], 4-mode, 5-mode and 7-mode models for both output projections were found to be stable. The model error for impulse response, however, decreases significantly if both modes corresponding to a traveling structure are included, as will be shown in Sec. V A (see Fig. 16). Balanced POD modes one, four, and five, and the corresponding adjoint modes for the five-mode output projection are shown in Fig. 14. Note that the structure of modes four and five in Fig. 14 is almost identical, except for a spatial phase shift of exactly one quarter of the periodic domain.

A. Impulse Response

Figure 15 (a) shows the perturbation energy growth as captured by three different standard POD models. It was observed that the inclusion of modes which come in pairs (see Fig. 12 (a)) in the basis used to form the reduced order models does not change the system behavior appreciably - the response of a model including modes 1–9 (not shown in figure) is virtually indistinguishable from the response of the model including only the first three modes. Hence, the traveling structure modes do not contribute significantly to the dynamics of this perturbation. The inclusion of the tenth mode, which is streamwise-constant, improves the performance significantly, and the model composed of only the first three modes and the tenth mode performs as well as one including the first ten modes. In the same fashion, including the mode pairs 11–12, 13–14 and 15–16 does not affect the model performance. Including the seventeenth mode, which is also a streamwise-constant mode, improves the performance further. The tenth and the seventeenth mode correspond to 0.025% and 0.0074% of the total energy. The low-order standard POD models were found to capture poorly the initial condition of the full simulation (this will be discussed in more detail in Sec. V D), so they were also started from different initial conditions at later times (before or around the peak energy growth), when the projection of the simulation onto POD modes is close to the full simulation data, and they still did not capture the correct peak and the subsequent decay of the energy.

On the other hand, the performance of very low-order BPOD models is significantly better. Figure 15 (b) shows the perturbation energy growth as captured by three different models. Although the two-mode model does not accurately capture the initial condition, it does represent the energy growth at later stages reasonably well. A three-mode BPOD model captures the kinetic energy of the full simulation very well except for the initial period. While more modes are needed to capture exactly the initial transient, if only the energy growth is of interest, the three-mode model is sufficient. This striking difference is an illustration of the advantage of balanced truncation - for standard POD it is difficult to know apriori which modes will be important for the system dynamics as demonstrated above and a good low-order model was found only after a careful examination of the mode basis which provided some insight into the underlying physics.

Figure 16 shows the error $\|G - G_r\|$ for BPOD using the two output projections and for standard POD. The error for the BPOD models quickly reaches the asymptotic values dictated by the output projection, although more modes are needed compared to the optimal perturbation case de-



(a) (b)

FIG. 15: (a) Three, four and five-mode POD models formed from the indicated modes (b) Two-mode and three-mode BPOD models. The very low-order BPOD models do not capture very well the initial transient, as shown in the inset. The BPOD models are from the ten-mode output projection.

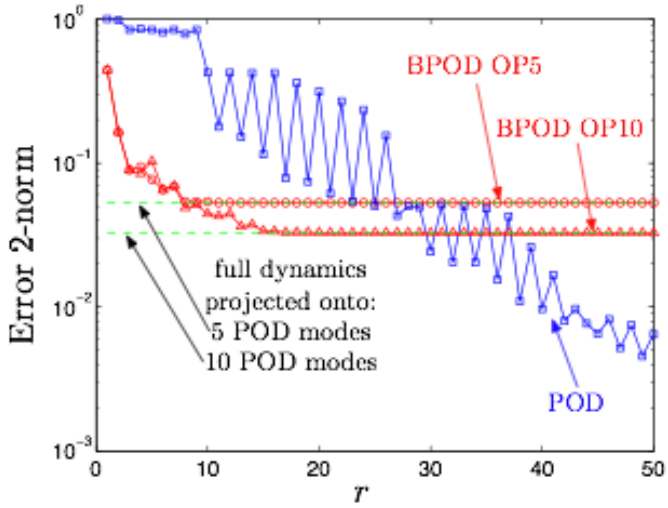


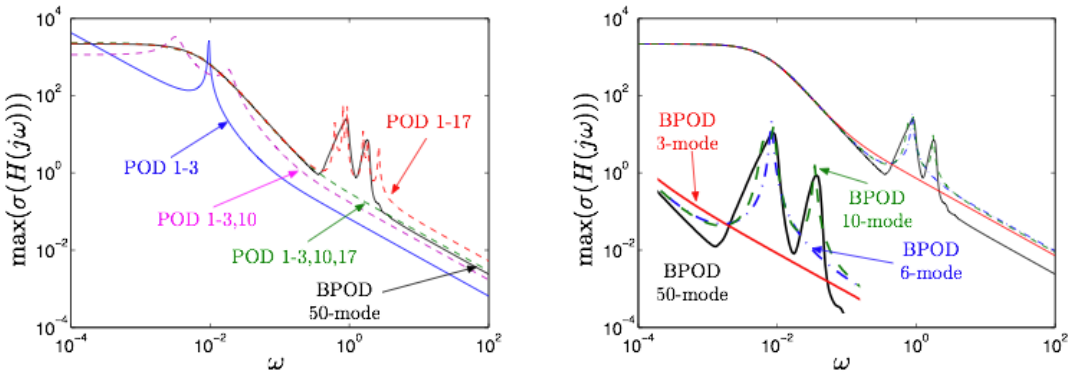
FIG. 16: Error 2-norms for localized actuator, showing POD models and BPOD at two output projections.

scribed in the previous section due to the more complex dynamics. Standard POD starts to match the performance of the ten-mode output projection BPOD only around rank 30, and also varies a lot with the model rank. This corresponds to the already observed fact that dynamically important POD modes are not highly ranked in terms of energy. Whenever the POD modes come in pairs, including only one of the modes results in deterioration of model performance. Eventually standard POD has better performance than BPOD, however recall that these POD models have r

outputs while the BPOD models have only s outputs. It should also be noted that some POD models exhibit sustained or very slowly decaying oscillations, and that the corresponding two-norms are in fact infinite. Since the simulation time for Fig. 16 is finite, the two-norms of such models appear to be large but finite as well. Although the four-mode BPOD models are stable for both output projections, including just one of the modes corresponding to a pair of equal HSVs 4–5 for the five-mode output projection deteriorates model performance, while we see a large decrease in the error when the fifth mode is included, as well as when we include subsequent pairs. For the ten-mode output projection, there are three equal HSVs 4–6, and a significant decrease in the error is seen only when we include all three of those modes (in particular, the error norm of the five-mode BPOD model is significantly larger than that of the four-mode model).

B. Frequency Response

Figure 17 shows the singular value Bode plots of standard POD and ten-mode output projection BPOD models for the localized disturbance. The frequency response of the 50-mode BPOD model, which is a very close approximation of the frequency response of the actual disturbance, has the shape of a low-pass filter with a break frequency of $0.01\ rad/s$ with two resonant peaks near $1\ rad/s$, which are similar to the peak observed for the single-wavenumber traveling structure perturbation in the previous section. We see that standard POD models again have spurious peaks at low model ranks. The addition of mode pairs corresponding to the traveling structures is necessary in order to reproduce the peaks around $1\ rad/s$ for both standard POD and BPOD, however BPOD captures those peaks with only the triple 4–6 and the mode pair 7–8, as well as modes 9 and 10 (Figure 17 (b)) while all standard POD modes 1–17 are needed to reproduce the same peaks and there are still spurious peaks. Since the peaks correspond to the low-energy traveling structures, it is not surprising that only a three-mode BPOD model performs so well in capturing the kinetic energy of the full simulation, as shown in the previous section. On the other hand, if the frequency response of the actuator around the frequency of $1\ rad/s$ needs to be captured accurately, the higher BPOD modes need to be included.

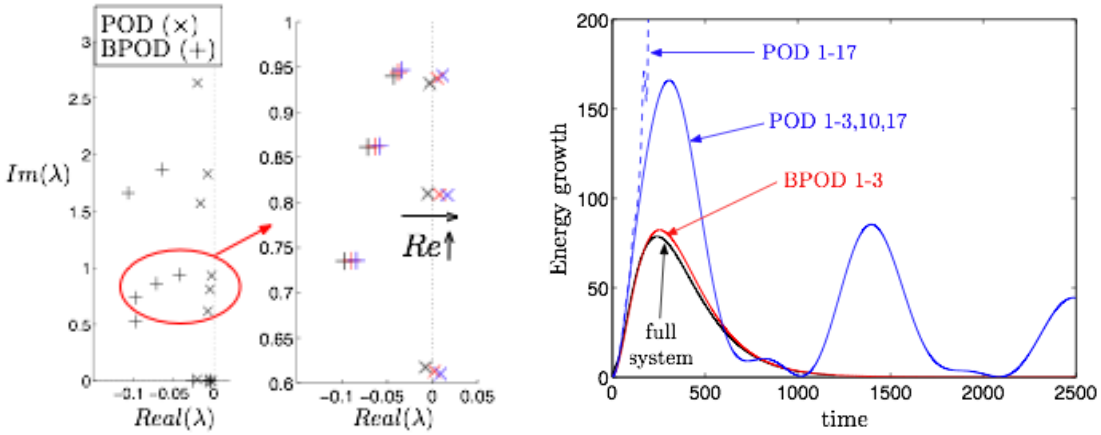


(a) (b)

FIG. 17: Singular value Bode plots for POD (a) and BPOD (b) models for the localized disturbance. The models are compared to a 50-mode BPOD model. The close-up in (b) shows that a six-mode BPOD model is needed to capture the larger resonant peak, and a ten-mode BPOD model captures both peaks.

C. Variation of Reynolds number

Figure 18 shows some of the eigenvalues of the 17-mode POD and BPOD models as the Reynolds number is increased. As in Section IV D, we use the modes from the design condition ($Re = 2000$ in this case) and form the models using Eq. (25). Both POD and BPOD models have eigenvalues on the real axis very close to the origin, which remain stable and correspond to the slow evolution of the streamwise-constant structures. At each Re, the eigenvalues of both models move towards the right half of the complex plane and while the BPOD model always remains stable, the standard POD model first appears marginally stable at $Re = 2500$ and then unstable at $Re = 3000$. The effect of the eigenvalues that move to the right half of the complex plane is clearly seen in Fig. 18 (b). A model that includes modes 1–17 grows unstable quickly at $Re = 3000$, showing that inclusion of modes which at design condition do not contribute signifcantly to the overall dynamics can significantly deteriorate the performance of the model at off-design condition. This can also be seen from the frequency responses shown in the previous section - even at design condition, the spurious high peaks correspond to marginally stable modes. Although stable, the 1,2,3,10,17 POD model is highly inaccurate at $Re = 3000$, with large peaks in the kinetic energy which decay very slowly, indicating the high sensitivity of those POD models which remain stable to a change in the Reynolds number. On the other hand, the three-mode BPOD model is still remarkably close to the full system.



(a) (b)

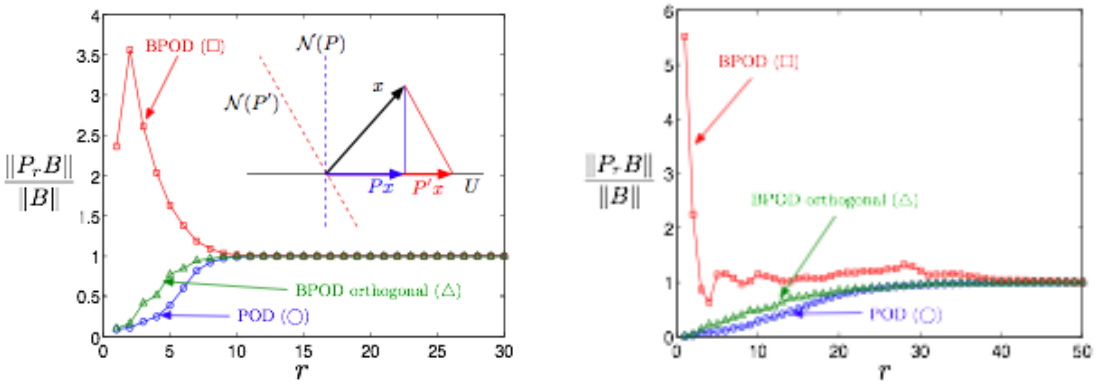
FIG. 18: (a) The eigenvalues of 17-mode models at $Re = 2000$, $Re = 2500$, and $Re = 3000$. (b) The performance of the models at the off-design condition of $Re = 3000$.

It is important to note here that as the Reynolds number is increased, the nonlinearity will have a stronger effect on the development of the disturbance and the reduced-order model may not be valid for higher Re in the first place. The comparison of the linear perturbation growth with a full nonlinear DNS solution is essential for a true validation of the models for controls applications, since we may be modeling the linearized flow well, but the linearized flow may not be a good approximation to the actual flow. This comparison is subject of current work.

D. Capturing of actuation

An important property of a reduced-order model is how well it captures the effects of the actuator in the original system, especially for models that are intended for developing controllers. In order for a reduced-order model to capture the effect of an actuator, it is necessary at a minimum for the input term in the equations (Bu in Eq. (4)) to be contained in the subspace used for projecting the equations. Note that here, even for the standard POD case, the effect of the "actuator" is partially included, since the dataset used for POD is generated by an impulsive input. One way to measure the degree to which the input "directions" are captured by the modes used in the model is to compute the projection of the columns of the input matrix B in Eq. (4) onto the basis modes. In the system we are considering here, B is a single column vector, representing the initial disturbance given to the system (or actuation via a body force in the center of the domain). Fig. 19 shows

the norm of the projection $\|P_r B\|/\|B\|$ of the standard and BPOD modes onto the input vector B, which is just the initial condition for each simulation. The balancing modes clearly capture the input direction with many fewer modes than standard POD: even very low-order models have a significant norm after projection, and in fact the norm of B after projection is almost always greater than the norm of B due to the non-orthogonal projection, as shown in Fig. 19. Any orthogonal projection such as P must satisfy $\|Px\| \leq \|x\|$, while for a non-orthogonal projection we may have $\|P'x\| > \|x\|$, which is the case for the first several BPOD modes. Clearly, the B matrix has a very small projection onto the POD modes for standard POD unless many modes are taken, so it is impossible for very low-order POD models to capture the response of an actuator without introducing more modes (such as the B matrix itself, Krylov subspaces, or shift modes [25]).



(a) (b)

FIG. 19: Norm of the projection of the B matrix (a single column vector) onto subspaces used for reduced-order models, (a) $\alpha = 1, \beta = 1$, (b) localized disturbance. The diagram in (a) illustrates the non-orthogonal projection used in BPOD.

E. Subspace comparison

The BPOD procedure uses both a different projection and a different set of modes in order to form reduced-order models, and we next look at a comparison of the two mode subspaces. A way to compare two subspaces is to compute $Tr(P_A P_B P_A) = T$, where P_A and P_B are the corresponding projection operators [5]. The trace of the matrix $Tr(P_A P_B P_A)$ as a function of the subspace rank is shown in Fig. 20 for the five-mode and ten-mode output projections where P_A and P_B are the orthogonal projectors onto the POD and BPOD subspaces respectively. The

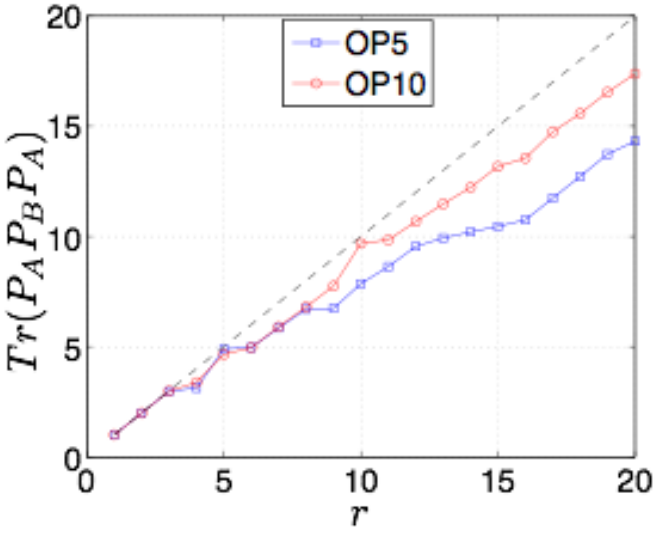


FIG. 20: Plots of $Tr(P_A P_B P_A)$ as a function of the subspace rank r for the localized perturbation for both output projections.

value of the trace T is the same as the subspace rank r at low order, indicating very similar modes (see Figs. 13 and 14). For both POD and five-mode OP BPOD, modes four and five are a pair of structures, and including both modes from the pair brings the value of s to almost exactly 5. It is interesting to observe that for the five-mode output projection, $r = T$ exactly at $r = 5$, while for the ten-mode output projection $r = T$ at $r = 10$, and that above those values the value of the trace is lower than the rank. This can be explained by the fact that BPOD is attempting to approximate the output projection of the data of the given rank. It is interesting to note that the subspaces including the first three POD and BPOD modes are virtually identical, indicating that the non-orthogonal Petrov-Galerkin projection via adjoint modes makes the enormous difference that we have seen in the performance of the corresponding models. As mentioned above, the POD basis is the basis of the most controllable modes, and is indeed optimal in capturing a given dataset, but as we have shown, it can fail to capture the dynamics correctly.

VI. CONCLUSIONS AND FURTHER WORK

We have shown that the balanced POD procedure offers several advantages over the standard POD technique for modeling linearized channel flow. Balanced POD makes the benefits of the systematic model reduction via balanced truncation available for large systems, resulting in mod-

els that capture the impulse response, frequency response and the effects of an actuator much better than their corresponding POD models. The differences between POD and BPOD are especially striking for localized actuators or disturbances which are often of interest in applications. As BPOD allows model reduction for very large systems, the localized perturbation models were obtained using three-dimensional fields in physical space, without the standard wavenumber decoupling which can be unwieldy for localized disturbances, since model reduction would have to be performed at each wavenumber pair resulting in models which are still of relatively high rank, and each wavenumber pair would need to be controlled separately. More importantly, this approach allows the extraction of dynamically dominant structures in physical space which may have contributions from many wavenumbers. This feature of the method will be very important in further applications of this method, and in particular in spatially developing flows.

There is an indication that the non-orthogonal projection onto the balancing modes with the use of adjoint modes plays a key role in the model performance. The subspaces spanned by the BPOD modes are very close to those spanned by the POD modes (which are optimal in representing the impulse response dataset). However, the dynamical models produced by BPOD are quite different, because of the non-orthogonal projection (or, equivalently, an orthogonal projection with a new inner product, defined by the observability Gramian [30]).

This work has focused on the linearized system in order to characterize the BPOD model reduction method. A true test of any controller is the application to the full nonlinear simulation, and verification of reduced perturbation growth, and this is a subject of our ongoing work. In addition, several techniques are available for obtaining nonlinear reduced-order models. One method for snapshot-based balancing for nonlinear systems has been introduced by Lall et al [18], but this involves considerably more computation than the present method, and is not feasible for large systems. A simpler, but presumably less effective approach is to project the full Navier-Stokes equations onto the BPOD modes computed for a linearized system; such a procedure involves no additional computational expense over the methods presented in this paper.

Although the body force actuation we discuss here would be desirable in applications, it is not as practically feasible as wall blowing and suction, which has been studied extensively and has shown promise both in computations and experiment [12, 21]. We also note that blowing and suction could be incorporated into the control framework described here in a relatively straightforward fashion, via the lifting approach described in Ref. [12]. Finally, the BPOD procedure is applicable to other flows that can be linearized, such as Couette flow, boundary layer flows, or flows over

airfoils. Moreover, the method can be applied to any large linear system and can therefore be a useful addition to the tools of modern control theory.

Acknowledgments

We gratefully acknowledge Dan Henningson, Shervin Bagheri, Espen Åkervik and Sunil Ahuja, as well as the anonymous referees, for helpful comments and suggestions. This work was supported by the National Science Foundation, grant CMS-0347239.

APPENDIX A: SNAPSHOT-BASED BALANCED TRUNCATION USING A CONTINUOUS ADJOINT

When computing the exact balanced truncation, the balancing transformation is found from the eigenvalue problem $W_c W_o T = T\Sigma^2$ where W_c and W_o are the controllability and observability Gramians with $A^+ = A^T$. We show here that, although the product $W_c W_o$ does not depend on the inner product on the state space used to define the adjoint system, the appropriate weight M needs to be included in the computation via the method of snapshots.

We can represent the weighted inner product of two vectors q_1 and q_2 as

$$\langle x_1, x_2 \rangle_M = \int_\Omega x_1^* M x_2 d\Omega \tag{A1}$$

where the domain of integration Ω is the Hilbert space itself. The star denotes the complex conjugate transpose. The inner product weight M is part of the definition of the Hilbert space itself. We define the so-called continuous adjoint of an operator A with respect to this inner product as

$$\langle Aq_1, q_2 \rangle_M = \left\langle q_1, A^+ q_2 \right\rangle_M \tag{A2}$$

We use the symbol $+$ in order to distinguish the adjoint from the standard matrix transpose A^T. From this definition it is easily shown that

$$\begin{aligned} \langle Ax, z \rangle_M = \left\langle x, A^+ z \right\rangle_M \quad &\Rightarrow \quad A^+ = M^{-1} A^T M \\ \langle Bu, x \rangle_M = \left\langle u, B^+ x \right\rangle \quad &\Rightarrow \quad B^+ = B^T M \\ \langle Cx, y \rangle = \left\langle x, C^+ y \right\rangle_M \quad &\Rightarrow \quad C^+ = M^{-1} C^T \end{aligned} \tag{A3}$$

In the above, we have assumed that the input and output spaces use the standard (unweighted) inner product. Next, we obtain for the Gramians:

$$\begin{aligned} G_c &= \int_0^\infty e^{At} B B^+ e^{A^+ t} dt = \int_0^\infty e^{At} B B^T M M^{-1} e^{A^T t} M dt \\ G_o &= \int_0^\infty e^{A^+ t} C^+ C e^{At} dt = \int_0^\infty M^{-1} e^{A^T t} M M^{-1} C^T C e^{At} dt \end{aligned} \tag{A4}$$

where G_c and G_o denote the Gramians obtained with the weighted inner product. Since the matrices M and M^{-1} are constant, we can take them out of the integrals, obtaining

$$G_c G_o = W_c W_o \tag{A5}$$

Thus, we have shown that balanced truncation does not depend on the choice of the inner product used to derive the adjoint system, and this allows us to use a convenient inner product. (In numerical simulations the 'simple' discrete adjoint $A^+ = A^T$ may in fact be more difficult to compute than a continuous adjoint which may retain a similar form of the equations; for instance, this is the case for linearized channel flow).

Next, we consider the computation of balancing and adjoint modes via the method of snapshots. From the definition of the empirical Gramians (Eq. (7)) it is easily shown that $Y^+ = Y^T M$ (recall that the snapshots of the adjoint simulations, which are the columns of Y, are given by $z(t) = e^{A^+ t} C^+$). Thus, we can write the SVD in Eq. (8) as

$$Y^T M X = U \Sigma V^T \tag{A6}$$

If we define the inverse of the balancing transformation as $\Psi^T = \Sigma^{-1/2} U^T Y^T$ we can easily compute the adjoint modes just from the SVD and from the adjoint snapshots. Recall that the *columns* of Ψ give the adjoint modes. The two sets of modes will now be bi-orthogonal with respect to the M inner product, so that $\Psi^T M \Phi = I$.

An alternative, more intuitive explanation is that, since both the direct and the adjoint snapshots 'live' in the state space, the correct inner product is that including the weight M (which is a part of the definition of the Hilbert space in which they reside). It is therefore this weighted inner product that should be used for forming the matrix for the SVD. Furthermore, the balancing and adjoint modes are bi-orthogonal with respect to this weighted inner product.

[1] E. Åkervik, J. Hœpffner, U. Ehrenstein, and D. S. Henningson. Optimal growth, model reduction and control in a separated boundary-layer flow using global eigenmodes. *J. Fluid Mech.*, 579:305–314, 2007.

[2] B. Bamieh and M. Dahleh. Energy amplification in channel flows with stochastic excitation. *Phys. Fluids*, 13(11):3258–3269, Nov. 2001.

[3] K. Butler and B. Farrell. Three-dimensional optimal perturbations in viscous shear flow. *Phys. Fluids A*, 4:1637–1650, 1992.

[4] G. E. Dullerud and F. Paganini. *A Course in Robust Control Theory: A Convex Approach*, volume 36 of *Texts in Applied Mathematics*. Springer-Verlag, 1999.

[5] R. Everson and L. Sirovich. Karhunen-Loève procedure for gappy data. *J. Opt. Soc. Am. A*, 12(8):1657–1664, Aug. 1995.

[6] B. Farrell. Optimal excitation of perturbations in viscous shear flow. *Phys. Fluids*, 31:2093–2102, 1988.

[7] B. F. Farrell and P. J. Ioannou. Accurate low-dimensional approximation of the linear dynamics of fluid flow. *J. Atmospheric Sci.*, 58:2771–2789, 2001.

[8] L. H. Gustavsson. Excitation of direct resonances in plane poiseuille flow. *Stud. Appl. Math.*, 75:227–248, 1986.

[9] A. Hanifi, P. J. Schmid, and D. S. Henningson. Transient growth in compressible boundary layer flow. *Phys. Fluids*, 8(3), March 1996.

[10] D. S. Henningson, A. Lundbladh, and A. V. Johansson. A mechanism for bypass transition from localized disturbances in wall-bounded shear flows. *J. Fluid Mech.*, 250:169–207, 1993.

[11] D. S. Henningson and P. J. Schmid. Vector eigenfunction expansions for plane channel flows. *Stud. Appl. Math.*, 87:15–43, 1992.

[12] M. Högberg, T. R. Bewley, and D. S. Henningson. Linear feedback control and estimation of transition in plane channel flow. *J. Fluid Mech.*, 481:149–175, 2003.

[13] P. Holmes, J. L. Lumley, and G. Berkooz. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, Cambridge, UK, 1996.

[14] M. Ilak and C. W. Rowley. Reduced-order modeling of channel flow using traveling POD and balanced POD. AIAA Paper 2006-3194, 3rd AIAA Flow Control Conference, June 2006.

[15] S. S. Joshi, J. L. Speyer, and J. Kim. A systems theory approach to the feedback stabilization of infinitesimal and finite-amplitude disturbances in plane poiseuille flow. *J. Fluid Mech.*, 332:157–184, 1997.

[16] M. R. Jovanović and B. Bamieh. Componentwise energy amplification in channel flows. *J. Fluid Mech.*, 534:145–183, 2005.

[17] J. Kim, P. Moin, and R. Moser. Turbulence statistics in fully-developed channel flow at low Reynolds number. *J. Fluid Mech.*, 177:133–166, Apr. 1987.

[18] S. Lall, J. E. Marsden, and S. Glavaški. A subspace approach to balanced truncation for model reduction of nonlinear control systems. *Int. J. Robust Nonlinear Control*, 12:519–535, 2002.

[19] A. J. Laub, M. T. Heath, C. C. Page, and R. C. Ward. Computation of system balancing transformations and other applications of simultaneous diagonalization algorithms. *IEEE Trans. Automat. Contr.*, 32:115–122, 1987.

[20] K. H. Lee, L. Cortelezzi, J. Kim, and J. Speyer. Application of reduced-order controller to turbulent flows for drag reduction. *Phys. Fluids*, 13(5):1321–1330, 2001.

[21] F. Lundell. Pulse-width modulated blowing/suction as a flow control actuator. *Experiments in Fluids*, 35:502–504, 2003.

[22] T. Min, S. M. Kang, J. L. Speyer, and J. Kim. Sustained sub-laminar drag in a fully developed channel flow. *J. Fluid Mech.*, 558:309–318, 2006.

[23] P. Moin and R. Moser. Characteristic-eddy decomposition of turbulence in a channel. *J. Fluid Mech.*, 200:471–509, 1989.

[24] B. C. Moore. Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE Trans. Automat. Contr.*, 26(1):17–32, Feb. 1981.

[25] B. Noack, K. Afanasiev, M. Morzyński, G. Tadmor, and F. Thiele. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.*, 497:335–363, 2003.

[26] G. Obinata and B. D. O. Anderson. *Model Reduction for Control System Design*. Springer-Verlag, 2000.

[27] S. A. Orszag. Accurate solution of the orr-sommerfeld stability equation. *J. Fluid Mech.*, 50:689–703, 1971.

[28] S. C. Reddy and D. S. Henningson. Energy growth in viscous channel flows. *J. Fluid Mech.*, 252:209–238, 1993.

[29] S. C. Reddy, P. J. Schmid, J. S. Baggett, and D. S. Henningson. On stability of streamwise streaks and

transition thresholds in plane channel flows. *J. Fluid Mech.*, 365:269–303, 1998.

[30] C. W. Rowley. Model reduction for fluids using balanced proper orthogonal decomposition. *Int. J. Bifurcation Chaos*, 15(3):997–1013, Mar. 2005.

[31] C. W. Rowley, T. Colonius, and R. M. Murray. Model reduction for compressible flows using POD and Galerkin projection. *Phys. D*, 189(1–2):115–129, Feb. 2004.

[32] C. W. Rowley and J. E. Marsden. Reconstruction equations and the Karhunen-Loève expansion for systems with symmetry. *Phys. D*, 142:1–19, Aug. 2000.

[33] P. J. Schmid and D. S. Henningson. *Stability and Transition in Shear Flows*, volume 142 of *Applied Mathematical Sciences*. Springer-Verlag, 2001.

[34] L. Sirovich. Turbulence and the dynamics of coherent structures, parts I–III. *Q. Appl. Math.*, XLV(3):561–590, Oct. 1987.

[35] T. R. Smith. *Low-dimensional models of plane Couette flow using the proper orthogonal decomposition*. PhD thesis, Princeton University, 2003.

[36] G. Tadmor. Observers and feedback control for a rotating vortex pair. *IEEE Trans. Contr. Sys. Tech.*, 12:36–51, 2004.

[37] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll. Hydrodynamic stability without eigenvalues. *Science*, 261:578–584, July 1993.

[38] K. Willcox and J. Peraire. Balanced model reduction via the proper orthogonal decomposition. *AIAA J.*, 40(11):2323–2330, Nov. 2002.

[39] K. Zhou and J. C. Doyle. *Essentials of Robust Control*. Prentice-Hall, 1998.